arXiv:1304.1942v1 [cs.NI] 6 Apr 2013

The effect of caching on a model of content and access provider revenues in information-centric networks[*]

F. Kocak and G. Kesidis
CS&E and EE Depts
Pennsylvania State University
University Park, PA, USA
{fwk5027,gik2}@psu.edu

T.-M. Pham and S. Fdida
CS Dept
UPMC
Paris, France
{tuan-minh.pham,serge.fdida}@lip6.fr

April 4, 2013

April 9, 2013

Abstract

In this paper, we consider a game between an Internet Service Provider (ISP, network access provider) and content provider (CP) on a platform of end-user demand. A price-convex demand-response is motivated based on the delay-sensitive applications that are expected to be subjected to the assumed usage-based prices (for priority service

---

[*] This research was supported by NSF CNS grant 1116626.

over best-effort which is the base service class for a flat-rate monthly fee). Both the Internet and proposed Information Centric Network (ICN) scenarios are considered, where the ICN case is different in the polarity of the side-payment (from ISP to CP in an ICN) and, more importantly here, in that content caching by the ISP is incented. A price-convex demand-response model is extended to account for content caching. The corresponding Nash equilibrium is derived and studied numerically.

1 Introduction

Network neutrality continues to be debated as its core economic issues as described in, *e.g.*, [10], have not been resolved. The debate concerns all participants in the enormous and growing Internet economy: Internet service (access) providers (ISPs), content providers (CPs, including providers of cloud computing services), end-user consumers, and government regulators.

In this paper, we consider a game between a single (or cooperating collective) CP and a single ISP on a platform of end-users served by both, *i.e.*, a two-sided market ([11, 13] developed games involving end-users and content providers on an ISP platform). We assume applications under consideration are delay-sensitive. monthly access costs for best-effort service. Applications requiring only best-effort service are not considered herein.

We are interested in the setting of the current Internet, as well as that of proposed Information-Centric Networks (ICNs), *e.g.*, [8, 16, 14]. In a related discussion, different networking scenarios for transit networks and content distribution networks (CDNs) were considered in [1], including those in which the CDN (CP) is incented to compensate the transit network (ISP) to cache its content. In this paper, the principle difference between the Internet and ICN settings is the direction of the side-payment between ISP and CP, similar to the difference between content-centric and access-centric networking as described in [16] (see Figures 5 and 6). Also, and more importantly as we do not model inter-provider competition herein (*e.g.*, [5, 1, 9]), the ISP is incentivized to cache in the ICN setting.

This paper is organized as follows. In Section 2, we summarize prior results on a simple ISP-CP game for the Internet setting. In Section 3, we adapt these results to the ICN setting and extend the model to account for content caching by the ISP. A key element of the extension is a price-convex demand-response motivated by delay-sensitivity of the applications/content under consideration. In Section 4, we give the results of a numerical study. Finally, we conclude in Section 5 with a short summary.

2 Background: The Internet model

2.1 A single ISP and CP with linear demand-response to price

Suppose there are two providers, one content (CP indexed 0) and the other access (ISP indexed 1), with *communal* consumer demand [6]. First suppose that the demand response to price is linear:

$$D = D_{\max} - d(p_1 + p_2), \qquad (1)$$

where d is demand sensitivity to the price, p_1 and p_2 are, respectively, the prices charged by the ISP and CP, and $D_{\max} > 0$ is the demand at zero usage based price[1]. Suppose the revenue of the ISP is

$$U_1 \quad = \quad (p_1 + p_s)D, \tag{2}$$

where p_s is the side payment from content to access provider. Similarly, the revenue of the CP is

$$U_2 \quad = \quad (p_2 - p_s)D. \tag{3}$$

Consider a noncooperative game played by the CP and ISP adjusting their prices, respectively p_2 and p_1, to maximize their respective revenues, with all other parameters fixed[2].

The following simple result was shown in [3, 5].

Theorem 1. *The interior Nash equilibrium*[3] *is*

$$p_1^* = \frac{D_{\max}}{3d} - p_s \quad \textit{and} \quad p_2^* = \frac{D_{\max}}{3d} + p_s$$

when

$$|p_s| \quad < \quad \frac{D_{\max}}{3d}, \tag{4}$$

---

[1] Note that ISPs are continuing to depart from pure flat-rate pricing (based on maximum access bandwidth) for unlimited monthly volume, *e.g.*, [15, 4].

[2] In particular, the fixed side-payment p_s is here assumed regulated. Note that the utilities are linear functions of p_s so that if p_s were under the control of one of the players, p_s would simply be set at an extremal value.

[3] In this paper, we do not consider boundary Nash equilibria, where at least one player is selecting an extremal value for one of their control parameters, often resulting in that player essentially opting out of the game, or maximally profiting from it at the expense of the other player. The boundary equilibria are also specified in [3].

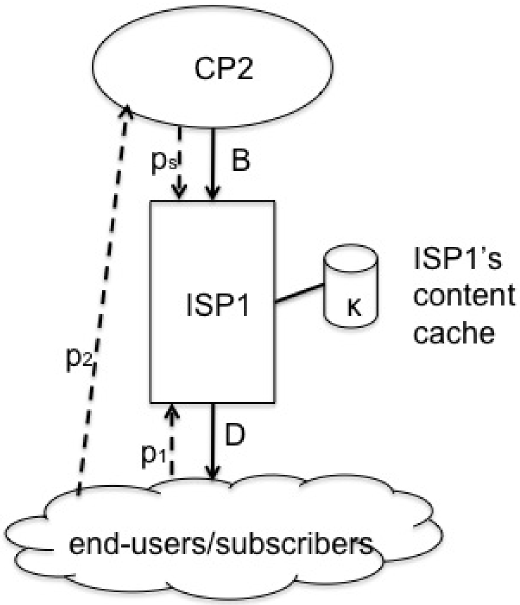

Figure 1: ISP and CP game on a platform of end-user demand-response

with player utilities

$$U_1^*, U_2^* = \frac{D_{\max}^2}{9d}.$$

Note that this result allows $p_s < 0$, *i.e.*, net side payment is from ISP to CP (remuneration for content instead of access bandwidth). But in the Internet setting, we take $p_s > 0$, whether there is direct side-payment from CP to ISP, or indirect payment through the peering contract [17, 7] between the residential ISP and the ISP of the CP (a contract that penalizes for asymmetric traffic exchange neutrally based on aggregate traffic volume).

In [5, 12], we showed that the the ISP may actually experience a reduction in revenue/utility with the introduction of side payments, using a communal demand model that had different demand-price sensitivity parameters d per provider type and also multiple providers of each type (*i.e.*, provider competition). Such a model was also considered in [2].

In [12], we used a convex, rather than linear, demand response to price, *e.g.*,

$$D \quad = \quad D_{\max}(1 - (p_1 + p_2)/p_{\max})^a, \tag{5}$$

where $a \geq 1$ and

$$p_{\max} \quad = \quad D_{\max}/d \;\; \text{when } a = 1.$$

This model was motivated in [12] by considering two different types of users, as follows. Suppose that (user-designated) premium class-of-service (CoS) applications are

- delay sensitive,
- given service priority by the ISP over best-effort applications for the bandwidth B available between CP and ISP,
- subjected to usage-based charges by the ISP at price p_1.

Best-effort applications exploit reserved-but-unused bandwidth ($\leq B$) by the premium CoS applications, and unreserved bandwidth if any. So, some delay-sensitive applications may be content with under best-effort CoS when demand for premium CoS is low (hence reserved-but-unused bandwidth is high). Thus, as demand increases for premium CoS applications, say because price $p = p_1 + p_2$ reduces, there may be *additional demand* owing to migration of delay-sensitive applications by more price-sensitive users who would otherwise tend to assign their delay-sensitive applications to best-effort CoS.

To better motivate this demand model, in this paper's Appendix A we derive a (more complex) price-convex demand response based on the delay-sensitivity of the usage-priced applications under consideration.

The following simple extension of Theorem 1 was shown in [12] by adding the first-order conditions $\partial U_i/\partial p_i = 0$, $i \in \{1, 2\}$, *cf.*, (7).

Theorem 2. *The interior Nash equilibrium for a strictly convex demand response D is*

$$p_1^* = p^* - p_s \quad \textit{and} \quad p_2^* = p^* + p_s, \tag{6}$$

where $p^* = p_1^* + p_2^*$ solves

$$2D(p^*) + p^* D'(p^*) \;=\; 0. \tag{7}$$

and $|p_s| < p^*/2$.

For the example of (5) with $a > 1$,

$$\begin{aligned} p^* &= \frac{2}{2+a} p_{\max}, & (8) \\ U_1^*, U_2^* &= \frac{p^*}{2} D(p^*) \;=\; \frac{D_{\max} p_{\max}}{2+a} \left(\frac{a}{2+a} \right)^a. & (9) \end{aligned}$$

Again, under communal demand response with only one provider of each type, neither $p^* = p_1^* + p_2^*$ nor U_1^* depend on p_s.

3 ICN model

In an ICN, residential users request content of the ISP, and the ISP (or some other intermediary resolver) decides the content provider, *e.g.*, [8]. Therefore in an ICN, it's reasonable to assume that the side-payment is from ISP to CP, *i.e.*, $p_s < 0$. Also, the ISP is motivated to cache content, unlike for our simple Internet case, to reduce the side payment (*i.e.*, avoid paying for, *e.g.*, the networking costs of the ISP-selected CP to transmit the user-requested content). Suppose that the ISP decides to cache a fraction κ of the content and this results in lower delay between the CP and ISP, (and a lower required side-payment to the CP, *cf.*, (11). If we model mean delay as $1/(B - D)$, where B is the service capacity between CP and ISP, then with caching factor κ, this delay is reduced to $1/(B - (1 - \kappa)D)$. For the model of Appendix B, the demand response:

- is increasing in caching factor κ,
- tends to convex in price as $\kappa \to 0$, and
- tends to linear in price as $\kappa \to 1$.

In the following for the ICN setting, we take a simplified form of demand response *with these properties*:

$$\begin{aligned} D &= D_{\max}(1 - (p_1 + p_2)/p_{\max})^{\kappa+(1-\kappa)a} \\ &= D_{\max}(1 - (p_1 + p_2)/p_{\max})^{a-\kappa(1-a)}. \end{aligned} \tag{10}$$

Also, note how both $D_{\max}$ and $p_{\max}$ are not affected by κ. Because ISP caching, the ISP and CP utilities generalize to Suppose the revenue of the ISP is

$$\begin{aligned} U_1 &= (p_1 + (1 - \kappa)p_s)D - c(\kappa), \qquad (11) \\ U_2 &= (p_2 - (1 - \kappa)p_s)D, \end{aligned}$$

again with $p_s < 0$, where $c(\kappa)$ is the cost of caching borne by the ISP.

We can use the results of Theorem 2 here, with parameters $(1-\kappa)p_s$ and $\kappa+(1-\kappa)a$ instead of p_s and a respectively, because the caching cost c component of U_1 does not depend on p_2 or p_1, and $|p_s| < p^*/2$ implies $|(1-a)p_s| < p^*/2$. We can conclude that the optimal utilities for ICN are

$$\begin{aligned} & U_1^* + c(\kappa), U_2^* \\ = & \frac{D_{\max} p_{\max}}{2+\kappa+(1-\kappa)a} \left(\frac{\kappa+(1-\kappa)a}{2+\kappa+(1-\kappa)a} \right)^{\kappa+(1-\kappa)a}. \end{aligned} \tag{12}$$

In the following section on numerical results, we consider performance at Nash equilibria as a function of κ (under the assumption that $|p_s| < p_{\max}/2$).

4 Numerical results

Figures 2-5 depict ISP utility $U_1^*/(D_{\max}p_{\max})$ with demand-exponent parameter $a = 2.0$ and caching cost of the form

$$c(\kappa) \quad = \quad bD_{\max}p_{\max}\kappa^n.$$

In Figure 2, $b = 0$ (*i.e.*, no cache cost, $c = 0$) and we see that U_1^* increases with caching factor κ. Figures 3 and 4 illustrate how linear cache cost ($n = 1$) leads to optimal $\kappa \in \{0, 1\}$: if $b \leq 0.04$ then optimal $\kappa = 0$, otherwise if $b \geq 0.05$ then optimal $\kappa = 1$. Figure 5 shows how the ISP utility may be concave in κ for quadratic cache cost ($n = 2$) - here for $b = 0.05$, optimal $\kappa \approx 0.4$. We can also compare the performance of ICN and Internet, where at $\kappa = 0$ for the Internet in these figures.

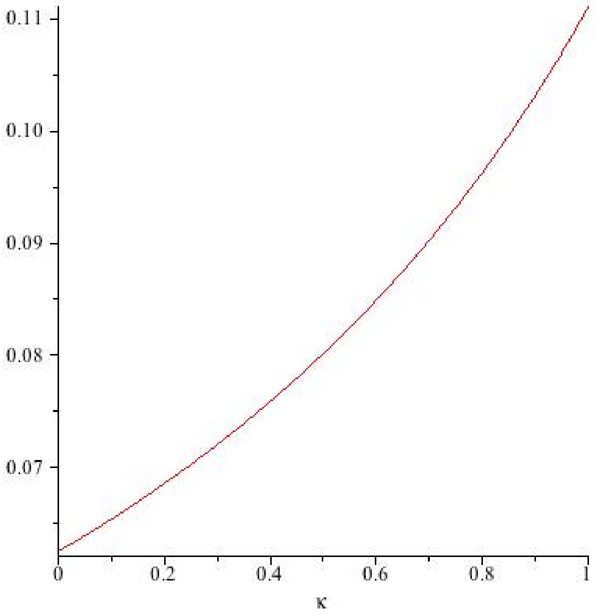

Figure 2: $U_1^*/(D_{\max}p_{\max})$ without caching cost

5 Summary

For an Information-Centric Network (ICN), we studied a two-player game involving a content provider and network-access provider (ISP) on a platform of price-convex end-user demand. In our ICN model, side-payments may flow from the ISP to the CP, *i.e.*, in the opposite direction than that of the current Internet. Also, depending on the cost of caching, the ISP may be incentivized to cache in the ICN setting, but there is no incentive for the ISP to cache in the Internet version of the considered game. For caching cost that is quadratic in caching factor, and price-quadratic demand-response, we showed how a fractional caching factor could be optimal at Nash equilibrium. In the future, we will extend this work to consider competition among like

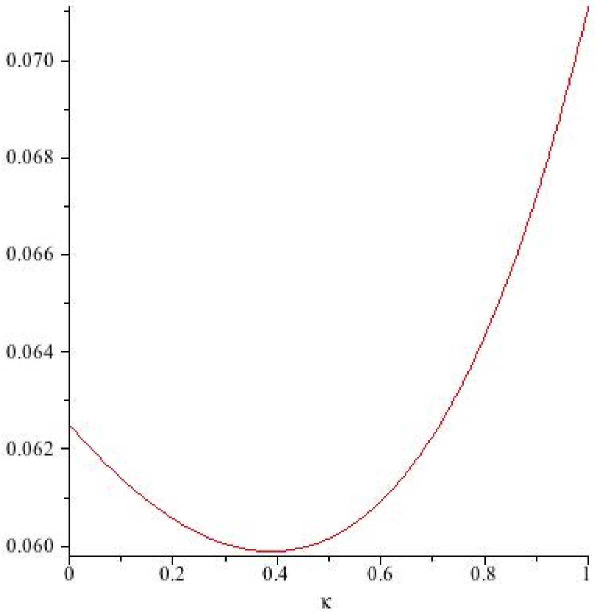

Figure 3: $U_1^*/(D_{\max} p_{\max})$ with linear caching cost, $b = 0.04$

providers (as in [5, 1, 9]), where side payments play a more significant role, and there may be coalitions among different providers, *i.e.*, eyeball ISPs.

References

[1] P. Agyapong and M. Sirbu. Economic Incentives in Content-Centric Networking: Implications for Protocol Design and Public Policy. In *Proc. 39th Telecommunications Policy Research Conference*, Arlington, Va, 2011.

[2] E. Altman, A. Legout, and Y. Xu. "Network non-neutrality debate: An economic analysis." In *Proc. IFIP Networking*, 2011.

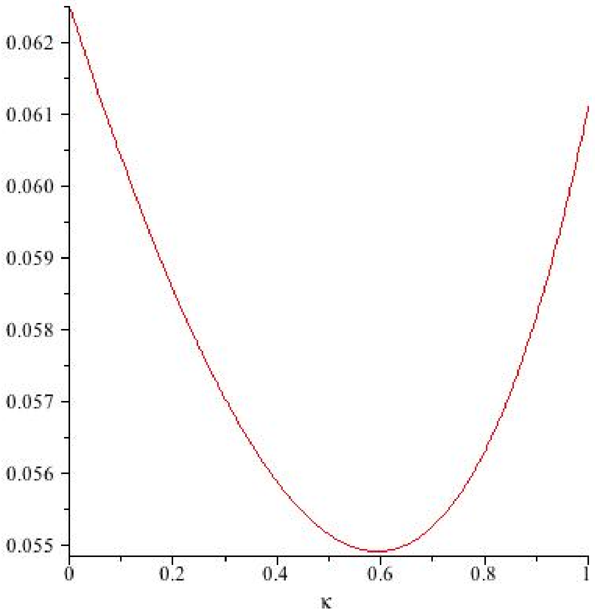

Figure 4: $U_1^*/(D_{\max}p_{\max})$ with linear caching cost, $b = 0.05$

[3] E. Altman, P. Bernhard, S. Caron, G. Kesidis, J. Rojas-Mora, S. Wong, "A study of non-neutral networks". *Telecommunication Systems Journal Special Issue on Socio-economic Issues of Next Generation Networks*, June 2011.

[4] K. Bode. AT&T To Impose Caps, Overages. Mar. 13, 2011. Available at: http://www.dslreports.com/shownews/Exclusive-ATT-To-Impose-Caps-Overages-113149

[5] S. Caron, G. Kesidis and E. Altman. "Application Neutrality and a Paradox of Side Payments", in *Proc. ACM ReArch*, Philadelphia, Nov. 30, 2010. See also http://arxiv.org/abs/1006.3894

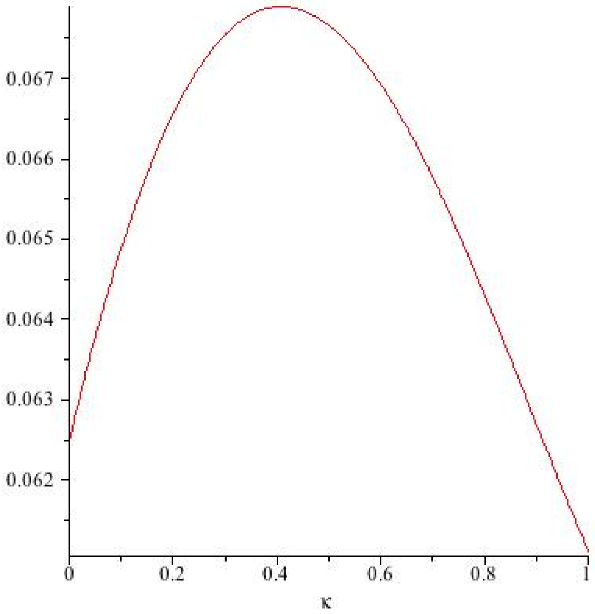

Figure 5: $U_1^*/(D_{\max}p_{\max})$ with quadratic caching cost, $b = 0.05$

[6] N. Economides. "Net Neutrality: Non-Discrimination and Digital Distribution of Content Through the Internet", *I/S: A Journal of Law and Policy*, **4**(2): pp. 209-233, 2008.

[7] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger and W. Lehr. Complexity of Internet Connections. In *Proc. 35th TPRC*, 2007.

[8] A. Ghodsi, T. Koponen, J. Rjahalme, P. Sarolahti, S. Shenker. Naming in Content-Oriented Architectures. In *Proc. ACM SIGCOMM ICN*, Toronto, Aug. 2011.

[9] F. Kocak, G. Kesidis, and S. Fdida. Network neutrality with content caching and its effect on access pricing. In *Smart Data Pricing*, S. Sen and M. Chiang (Eds.), Wiley, 2013.

[10] R. Hahn and S. Wallsten, “The Economics of Net Neutrality”, *Economists’ Voice*, **3**(6), pp. 1-7, 2006.

[11] P. Hande, M. Chiang, R. Calderbank, and S. Rangan, “Network pricing and rate allocation with content provider participation”, *Proc. IEEE INFOCOM*, 2009.

[12] G. Kesidis. Side-payment profitability under convex demand-response modeling congestion-sensitive applications. In *Proc. IEEE ICC*, Ottawa, Canada, June 2012.

[13] J. Musacchio, G. Schwartz and J. Walrand, “A Two-Sided Market Analysis of Provider Investment Incentives With an Application to the Net-Neutrality Issue”, *Review of Network Economics*, **8**(1), 2009.

[14] T.-M. Pham, S. Fdida, and P. Antoniadis. Pricing in Information-Centric Network Interconnection. Technical report, University Pierre et Marie Curie (UPMC), Paris, January 2013.

[15] A. Shin. Who’s the bandwidth bandit? *The Washington Post*, Oct. 4, 2006. Available at: http://blog.washingtonpost.com/thecheckout/2006/10/bandwidth_bandit.html

[16] D. Trossen and A. Kostopoulos. Techno-Economic Aspects of Information-Centric Networking. *Journal of Information Policy* **2**, p. 26-50, 2012.

[17] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V.V. Vazirani. How Many Tiers? Pricing in the Internet Transit Market. In *Proc. ACM SIGCOMM*, 2011.

[18] R.W. Wolff. *Stochastic Modeling and the Theory of Queues*. Prentice-Hall, Englewood Cliffs, NJ, 1989.

Appendix A: Explanation of convex demand response

We "implicitly" model the demand $D = [g(D)]^+$ with

$$g(D) = (D_{\max} - dp)\left(1 - \frac{\lambda}{B - D}\right) \Big/ \left(1 - \frac{\lambda}{B}\right), \tag{13}$$

where

- B is the bandwidth reserved *between CP and ISP* for delay sensitive applications paying usage-based prices,
- λ is demand sensitivity to mean delay, here modeled as $1/(B - D)$ (an expression for mean delay taken from the M/M/1 queue [18]).

Here, $\lambda > B - D$ results in zero demand D. That is,

$$D = [g(D)]^+.$$

Letting

$$\tilde{D} := (D_{\max} - dp)/(1 - \lambda/B) = D_{\max}(1 - p/p_{\max})/(1 - \lambda/B),$$

and assuming

$$\tilde{D} > 0,$$

we can find the interior fixed-point D of g^+ (*i.e.*, fixed point of g), giving the "explicit" demand response

$$D = \tfrac{1}{2}\left[(B+\tilde{D}) - \sqrt{(B-\tilde{D})^2 + 4\lambda\tilde{D}}\right]. \tag{14}$$

It's easy to see that this demand response has the following intuitive properties:

- $D \to D_{\max}$ as $B \to \infty$ and $p \to 0$
- D is a convex function of $\tilde{D}$ when $B > \lambda$, and hence also a convex function of price p (as assumed in [12, 9]).

There are obviously many alternative demand models with similar properties.

Appendix B: Explanation of convex demand response, increasing in caching factor

As a result of ISP caching, only a fraction $(1-\kappa)$ of the demand D is transmitted through the the bandwidth B between ISP and CP. So, (13) is modified to

$$\begin{aligned} g_\kappa(D) &= (D_{\max} - dp)\left(1 - \frac{\lambda}{B-(1-\kappa)D}\right) / \left(1 - \frac{\lambda}{B}\right) \\ &= (D_{\max} - dp)\left(1 - \frac{\lambda/(1-\kappa)}{B/(1-\kappa) - D}\right) / \left(1 - \frac{\lambda/(1-\kappa)}{B/(1-\kappa)}\right) \end{aligned}$$

So, solving $D = g_\kappa(D)$ results in (14) with B and λ replaced by $B/(1-\kappa)$ and $\lambda/(1-\kappa)$, respectively:

$$D = \tfrac{1}{2}\left[(\frac{B}{1-\kappa} + \tilde{D}) - \sqrt{(\frac{B}{1-\kappa} - \tilde{D})^2 + 4\frac{\lambda}{1-\kappa}\tilde{D}}\right]. \tag{15}$$

So, as $\kappa \to 0$, the demand tends to (14), *i.e.*, convex in price. On the other hand, as $\kappa \to 1$, the demand tends to linear in price (1).

Since, $g_\kappa(D) = g_0((1-\kappa)D) := g((1-\kappa)D)$, is decreasing in $(1-\kappa)D$ (hence increasing in caching factor κ), the solution

$$D_\kappa \ = \ g_\kappa(D_\kappa)$$

is an increasing function of caching factor κ (in particular) $D_\kappa \geq D_0$). To see this, note that

$$D_0 \ = \ g_0(D_0) \ < \ g_0((1-\kappa)D_0) \ = \ g_\kappa(D_0).$$

So, if $D_\kappa \leq D_0$, then we would have

$$D_\kappa \leq D_0 \ < \ g_\kappa(D_0) \ \leq \ g_\kappa(D_\kappa),$$

which contradicts the definition of D_κ in the first display above.